Database Administration

July 2006

Topics

- **This course addresses:**
  - **System design**
  - **DBA responsibilities**
  - **Starting and stopping servers**
  - **Creating database devices and logical volumes**
  - **Installing databases and patches**
  - **Configuring databases**
  - **Working with indexes, segments, and caches**
  - **Establishing database security**
  - **Copying, replicating, and extracting data**
  - **Replication system administration**
  - **Performance monitoring, tuning, and problem reporting**
  - **Ensuring database quality**
  - **Generating reports**
  - **Sybase Troubleshooting**
  - **Oracle Procedures**

System Overview
General Design

- **The system is designed to:**
  - **Receive data from external sources**
  - **Save those data in either long-term or permanent storage**
  - **Produce higher-level data products from the received data**
  - **Support access to the data by scientists and other registered clients**

System Overview
Information Model

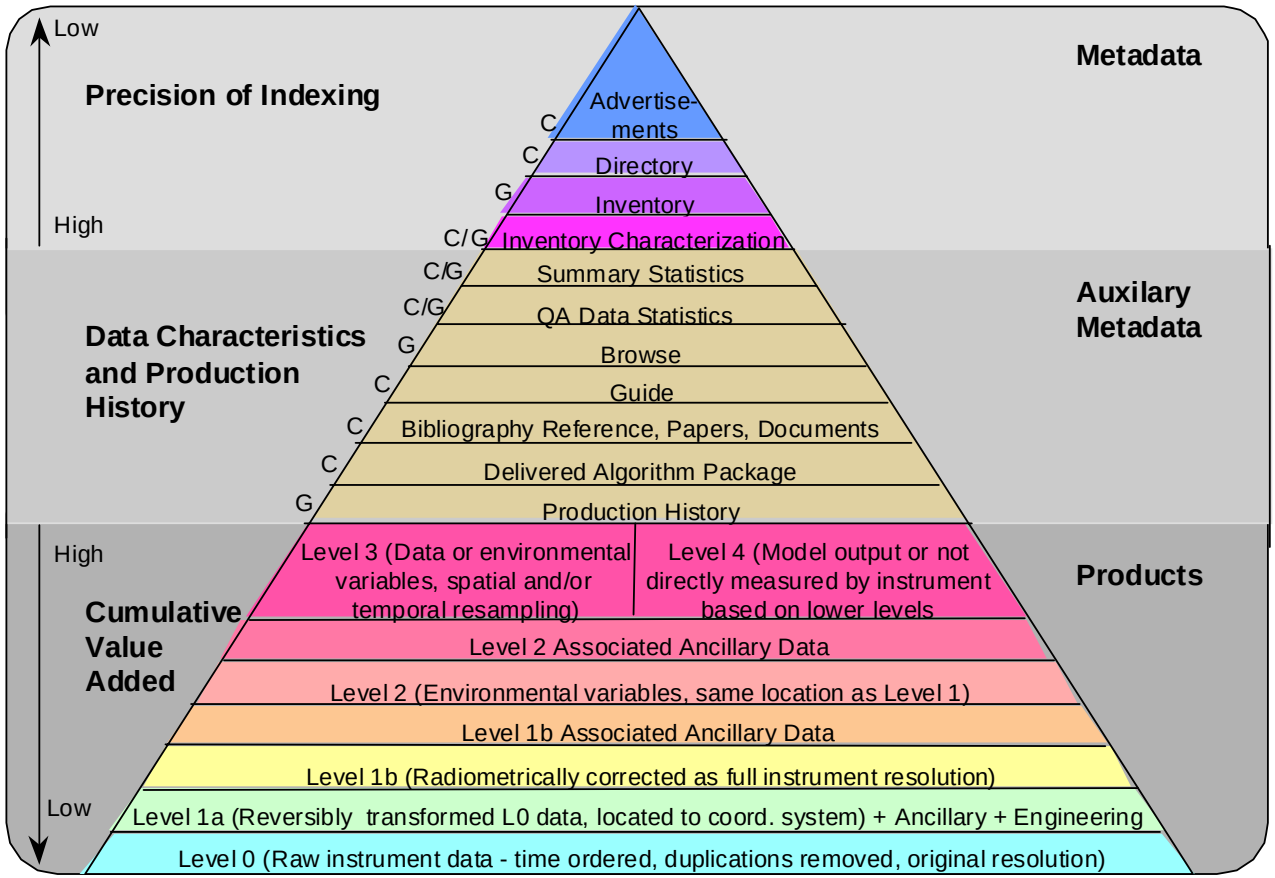

System Overview
Data Products

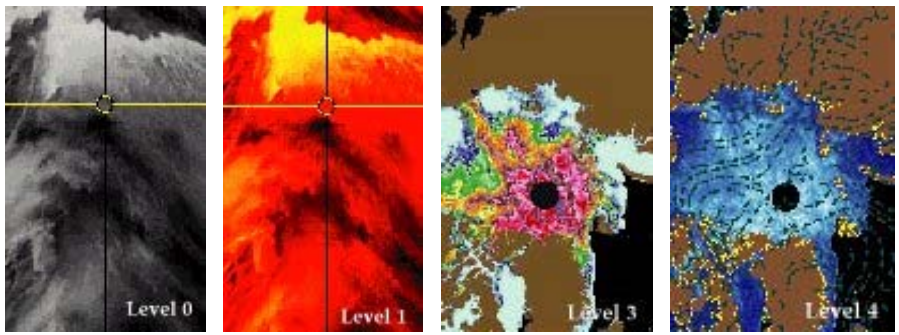

Level	Definition
0	Reconstructed, unprocessed instrument/payload data at full resolution; any and all communications artifacts (e.g., synchronization frames, communications headers, duplicate data removed)
1A	Reconstructed unprocessed instrument data at full resolution, time-referenced, and annotated with ancillary information, including radiometric and geometric calibration coefficients and geo-referencing parameters (e.g., platform ephemeris computed and appended but no applied to the Level 0 data)
1B	Level 1A data that have been processed to sensor units (not all instruments will have a Level 1B equivalent)
2	Derived geophysical variables at the same resolution and location as the Level 1 source data
3	Derived geophysical variables mapped on uniform space-time grid scales, usually with some completeness and consistency
4	Model output or results from analyses of lower-level data (e.g., variables derived from multiple measurements)

System Overview
Context Diagram

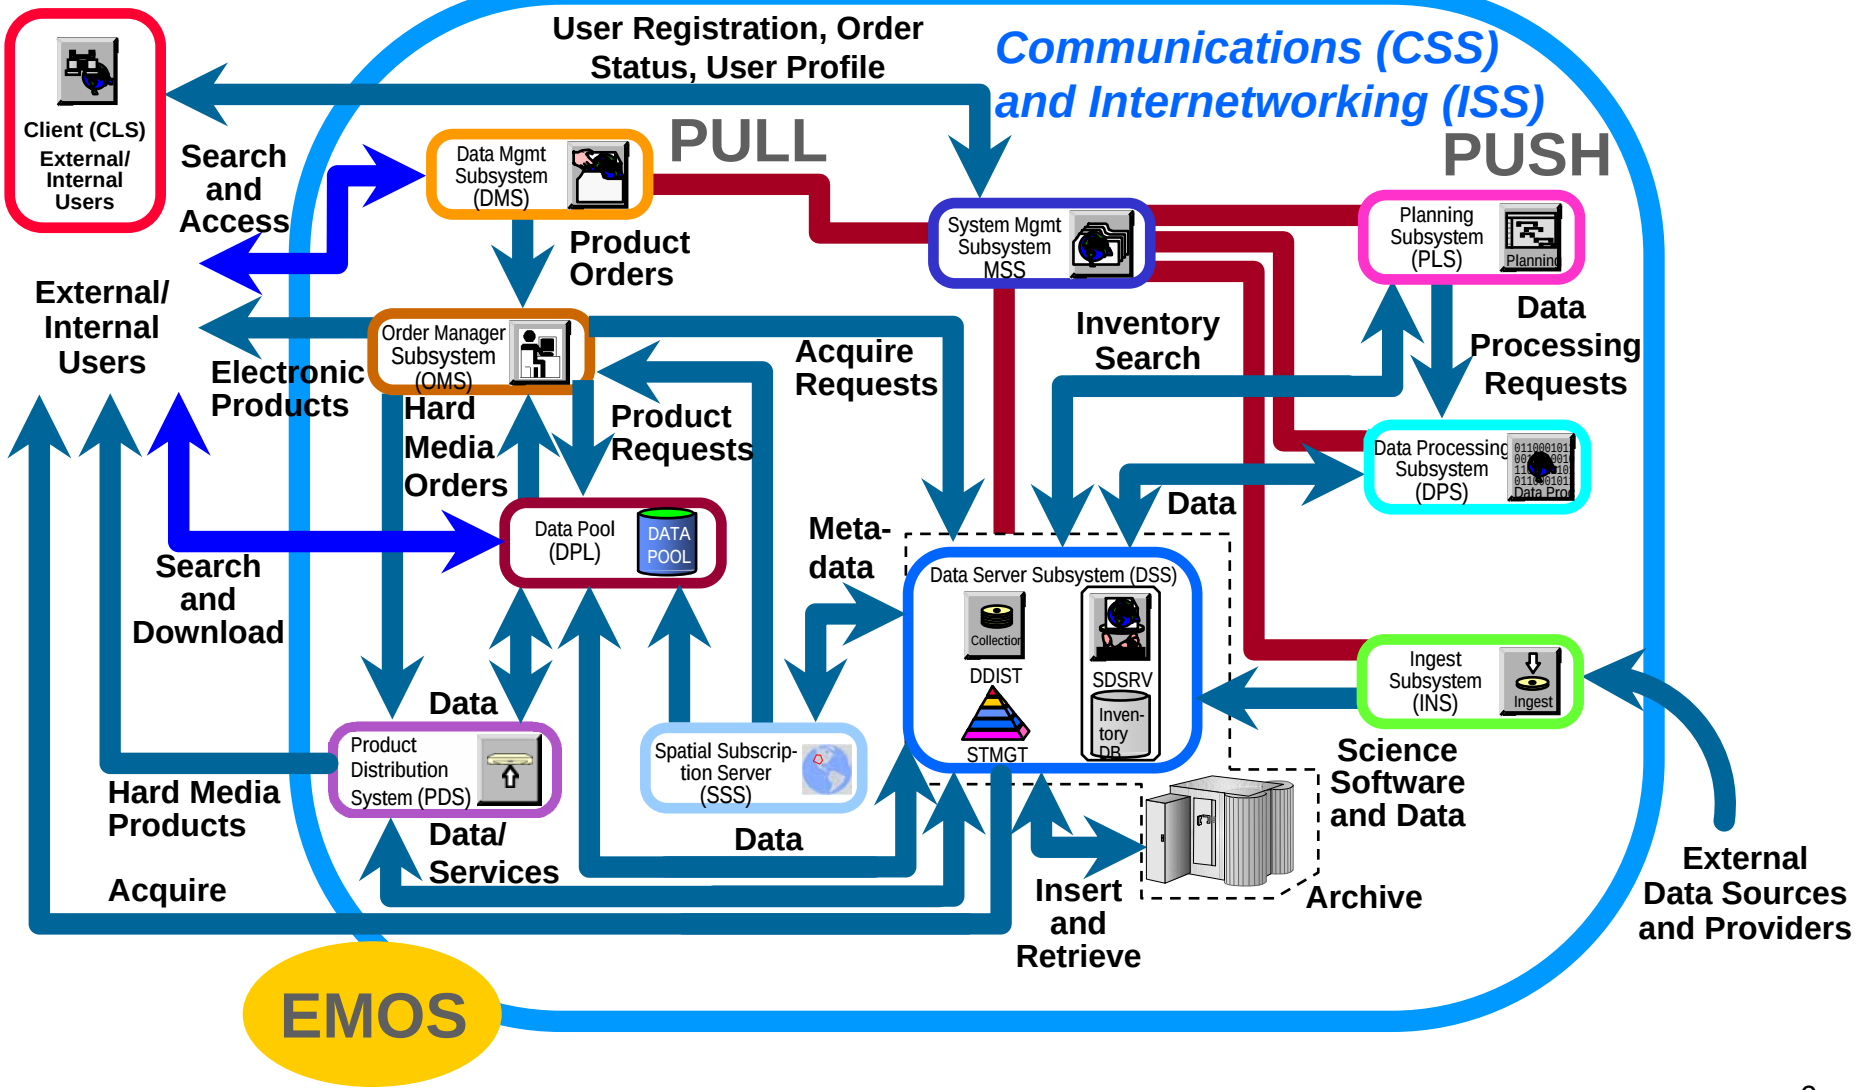

System Overview
Custom Databases

Database Name	Document Number	DB Software	No. of Tables	Logical Categories
Science Data Server Subsystem (SDSRV)	311-EMD-004	Sybase	154	Database Version Information
				System Management Data
				Collection, Granule Metadata
				DAP Metadata
				Spatial Metadata
				Data Originator Metadata
				Granule Metadata
				Contact Metadata
				Collection Metadata
				Temporal Metadata
Planning and Data Processing Subsystems (PDPS)	311-EMD-003	Sybase	80	Database Version Information
				Planning Data
				Data Processing Data
Data Management Subsystem (DMS)	311-EMD-001	Sybase	63	Database Versioning
				Attribute/Term Definitions
				Collection Metadata
				Information Management
Storage Management and Database Distribution Subsystems (STMGT & DDIST)	311-EMD-005	Sybase	67	Database Version Information
				Data Distribution
				Archive Services
				Request Handling
				Server Configuration
				Cache Management
				Media Operations
				FTP Services
				Staging Disk Operations
				GR Cleanup
Ingest Subsystem (INS)	311-EMD-002	Sybase	25	Database Version Information
				Datatype Information
				Configuration Data
				Active Requests
				Validation Data
				Table Locking Information

Database Name	Document Number	DB Software	No. of Tables	Logical Categories
Registry (REGIST)	311-EMD-008	Sybase	12	Database Version Information
				Security Information
				Registered Parameter Info
Systems Management Subsystem (MSS)	311-EMD-007	Sybase	19	Database Version Information
				Order Information
				Site Information
				Validation Data
				User Data
Subscription Server (SUBSRV)	311-EMD-006	Sybase	8	Database Version Information
				Subscription Information
				Event Information
NameServer (NM)	311-EMD-010	Sybase	2	Database Versioning
				NameServer
Product Distribution System (PDS)	311-EMD-009	Oracle	28	PDS Interface Server Order Data
				PDS Job Data
Data Pool (DPL)	311-EMD-013	Sybase	67	Collection Metadata
				Granule Metadata
				Insert Action Data
Order Manager Server (OMS)	311-EMD-011	Sybase	31	Queue/Status Information
				Request Information
				Intervention Information
Spatial Subscription Server (SSS)	311-EMD-012	Sybase	39	Database Version Information
				Subscription Information
				Event Information
				Action Information

System Overview
COTS Databases

Subsystem	COTS Product/Database Name	DB Software
PDPS	AutoSys	Sybase
MSS	Remedy	Sybase

System Overview
Flat Files

Flat Files

Database	Flat File Attributes			
	Usage	**Types**	**Formats**	**Descriptions**
SDSRV	Yes	UNIX flat file; ELF 32-bit MSB dynamic lib SPARC Version 1, dynamically linked	Variable length, Dynamic Link Library (DLL)	Log files, configuration files, template used to validate ESDTs on installation, uniquely named ESDT file descriptors, generic to ESDT-specific processing capabilities
PDPS	Yes	Text	ODL	Science metadata ODL file template
DMS	Yes	UNIX flat file	Variable length	Log files, configuration files
STMGT & DDIST	Yes	UNIX flat file	Variable length	Disk index files, staging data information, resource lists
INS	Yes	UNIX flat file	Variable length	Log files, configuration files, data delivery records
REGIST	No			
MSS	Yes	ASCII, binary	Single line records, one/two fields; EcAgEvent objects; MsAgMgmtHandle object; integers; string lists	Accountability component files, subagent component files
SUBSERV	No			
NM	No			
PDS	Yes	Text	ODL, Variable length	Production parameter files, status files, order data
DPL	Yes	ASCII	Variable length	For Data Pool Access Statistics Utility, temporary storage of data to be exported to database
OMS	No			
SSS	No			

System Overview
Resident Databases

Resident Databases

Databases	SMC	DAAC			
		GSFC	LP DAAC	LaRC	NSIDC
Custom					
Science Data Server Subsystem (SDSRV)		✔	✔	✔	✔
Planning and Data Processing Subsystem (PDPS)		✔	✔	✔	
Data Management Subsystem (DMS)	✔	✔	✔	✔	✔
Storage Management and Database Distribution Subsystems (STMGT & DDIST)		✔	✔	✔	✔
Ingest Subsystem (INS)		✔	✔	✔	✔
Registry (REGIST [MSS])	✔	✔	✔	✔	✔
Systems Management Subsystem (MSS)	✔	✔	✔	✔	✔
Subscription Server (SUBSRV)		✔	✔	✔	✔
NameServer (NM)	✔	✔	✔	✔	✔
Product Distribution System (PDS)		✔	✔	✔	✔
Data Pool (DPL)		✔	✔	✔	✔
Order Manager Subsystem (OMS)		✔	✔	✔	✔
Spatial Subscription Server (SSS)		✔	✔	✔	✔
Replication Server System Database (RSSD)		✔	✔	✔	✔
COTS					
AutoSys (PDPS)	✔	✔	✔	✔	
Remedy (MSS)	✔	✔	✔	✔	✔

System Overview
Database Replication

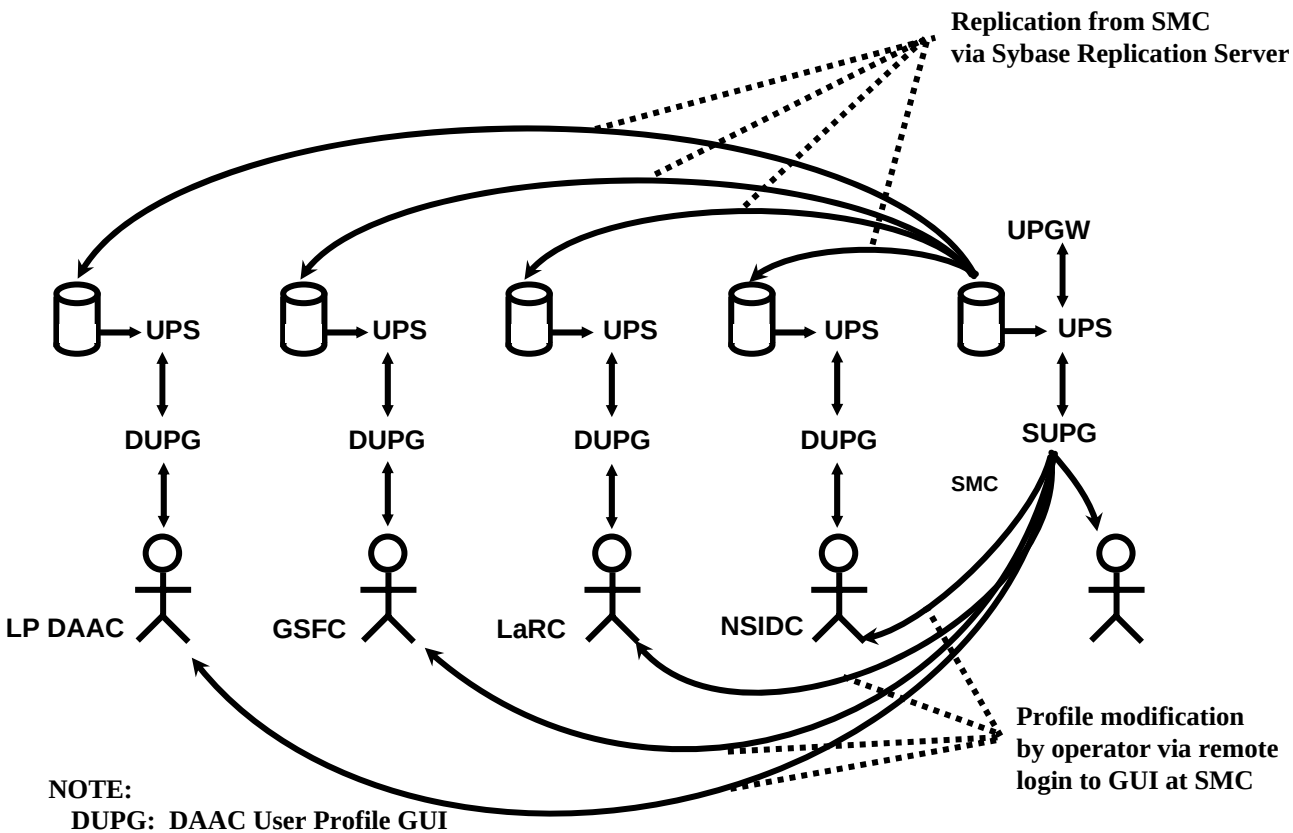

NOTE:
- DUPG: DAAC User Profile GUI
- SUPG: SMC User Profile GUI
- UPS: User Profile Server
- UPGW: User Profile Gateway

System Overview
Location of Principal Database Components

Name	Variant	Vendor	Principal Directory	Comments
Software Developer's Kit (formerly Open Client)	PC	Sybase	c:\windows\system	
Software Developer's Kit (formerly Open Client)	SGI	Sybase	/tools/sybOCv12.5.1	Just utilities, not libraries
Software Developer's Kit (formerly Open Client)		Sybase	/tools/sybOCv12.5.1	
Oracle Developer	SGI	Oracle	TBI	PDS only
Oracle Enterprise 8I	SGI	Oracle	TBI	PDS only/Oracle Forms 4.5 bundled
Replication Server	SUN	Sybase	/usr/ecs/OPS/COTS/sybase1151	
Replication Server Manager	SUN	Sybase	/usr/ecs/OPS/COTS/sybase1151	
ASE Server Monitor Client/Svr	SUN, SGI	Sybase	/usr/ecs/<mode>/COTS/sybase	At DAAC discretion/ required for launch
Spatial Query Server (SQS)	SGI	Autometrics	/usr/ecs/OPS/COTS/sqs_322/bin	
Sybase Adaptive Server	SUN, SGI		/usr/ecs/OPS/COTS/sybase1151	/usr/ecs/OPS/COTS/sybase _1151 is an acceptable install dir.
Sybase Adaptive Svr Enterprise	SUN, SGI	Sybase	/usr/ecs/OPS/COTS/sybase1151	/usr/ecs/OPS/COTS/sybase _1151 is an acceptable install dir.
Sybase Adaptive Svr Enterprise	SGI	Sybase	/usr/ecs/OPS/COTS/sybase_1193	
Sybase Adaptive Svr Enterprise	Sun	Sybase	/usr/ecs/OPS/COTS/sybase1151	/usr/ecs/OPS/COTS/sybase _1151 is an acceptable install dir.
Sybase Central	PC	Sybase	TBI	

System Overview
Database Management Implementation

- **System databases are primarily based on Sybase software. Only PDS uses Oracle software. Primary components include:**
  - **Sybase Adaptive Server Enterprise (ASE)**
  - **Other Sybase Components:**
    - **Spatial Query Server (SQS)**
    - **Replication Server (RS)**
  - **Oracle Enterprise**

System Overview
Sybase ASE Components

Type	Component	Description	Sub-Components and Features
Client	Sybase Central	A Windows application for managing Sybase databases. Helps manage database objects and perform common administrative tasks.	Connecting to, disconnecting from, and stopping servers
			Troubleshooting Adaptive Server problems
			Managing data caches
			Managing Adaptive Server physical resources
			Creating, deleting, backing up, and restoring databases
			Creating and deleting Adaptive Server logins, creating and deleting database users and user groups, administering Sybase roles, and managing object and command permissions
			Monitoring Adaptive Server performance data and tuning performance parameters
	Sybase Central Plug-Ins	Each server product is managed by a service provider plug-ins that coexists with other service providers in the Sybase Central framework.	ASE Plug-In
			SQS Plug-In
			Replication Server Manager (RSM). Provides the ability to manage, monitor, and troubleshoot most replication system components (primary and replicate database servers, Replication Servers, Replication Agents, and database gateways).
	Software Developer's Kit (formerly Open Client)		CS-Library, which contains a collection of utility routines used by all client applications.
			Client-Library and DB-Library, which contain a collection of routines that are specific to the programming language being used in an application
			Net-Library, which contains network protocol services that support connections between client applications and Adaptive Server.
			Utilities: isql – an interactive query processor that sends commands to the RDBMS from the command line. bcp – a program that copies data from a database to an operating system file, and vice versa. defncopy - a program that copies definitions of database objects that from a database to an operating system file and vice-versa.

Type	Component	Description	Sub-Components and Features
Server	Adaptive Server (ASE)	Sybase's high-performance RDBMS	
	Backup Server(TM)	A server application that runs concurrently with Adaptive Server to perform high-speed on-line database dumps and loads.	
	Adaptive Server Monitor	Monitor Server	Allows capture, display, and evaluation of Adaptive Server performance data and tune Adaptive Server performance
		Historical Server	Writes the data to files for offline analysis

System Overview
Sybase Central

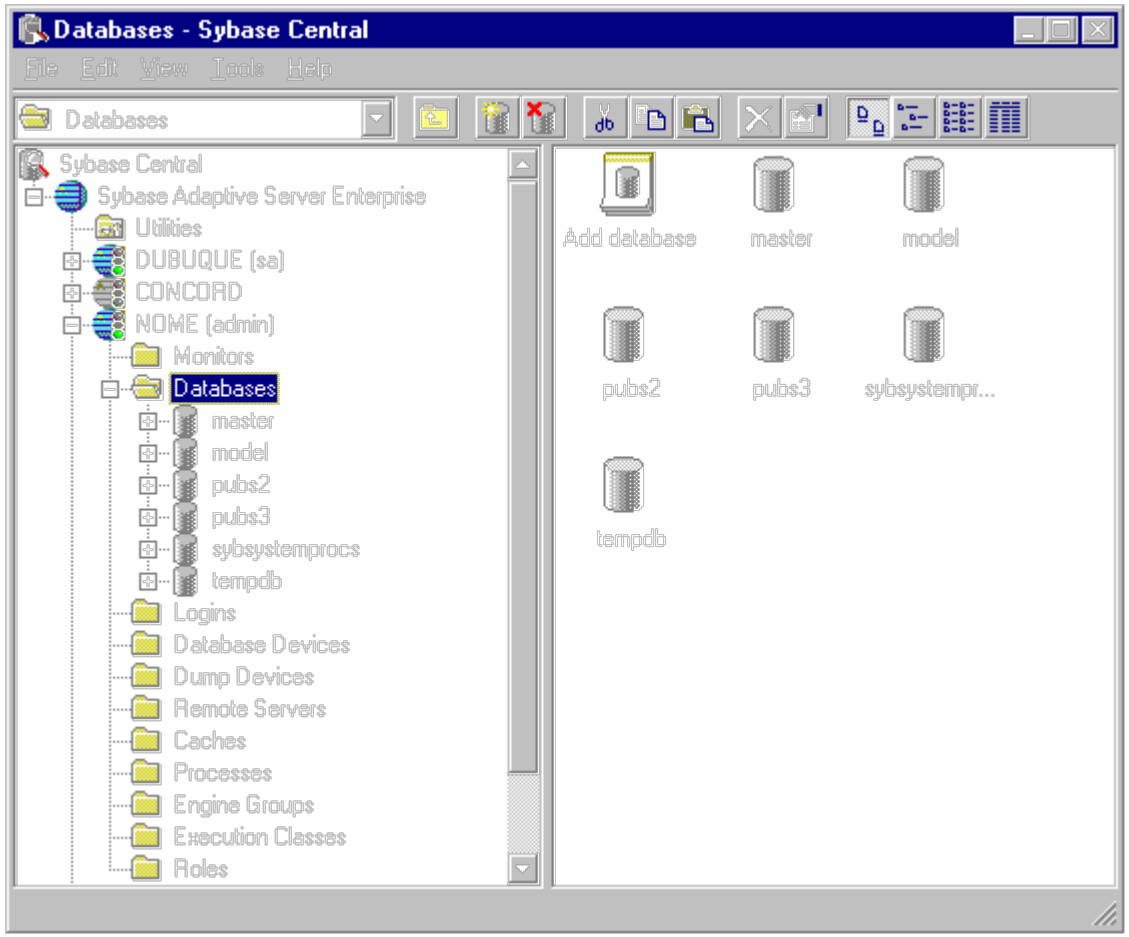

System Overview
Hardware, Software, and Database Mapping

- **Baseline information available at:**

  http://cmdm-ldo.raytheon.com/baseline/

- **Link to Technical Documents**
  - **920-TD*x*-001: Hardware-Design Diagram**
  - **920-TD*x*-002: Hardware-Software Map**
  - **920-TD*x*-009: DAAC HW Database Mapping**

DBA Responsibilities
Basic Responsibilities

- **Performing the database administration utilities**
  - **Such as database backup, maintenance of database transaction logs, and database recovery**
- **Monitoring and tuning the database system (e.g., the physical allocation of database resources)**
- **Maintaining user accounts for the users from the external system**
  - **Creating user registration and account access control permissions in the security databases**
- **Creating standard and ad hoc security management reports**
- **Working with EMD sustaining engineering and DAAC system test engineers to set up a test environment as needed**

DBA Responsibilities
Basic Responsibilities (Cont.)

- **Working with the data specialist on information management tasks involving databases, data sets, and metadata management**
- **Consolidating event reports into a site event history database for reporting activities to the SMC on a regular basis**
- **Performing daily database synchronization**
- **Administering the Replication Server System Database (RSSD)**

DBA Responsibilities
Routine Tasks

Time Period	Task	Importance	Found In …
Daily	Capture database configurations	Absolutely necessary for database recovery if problems occur	Configuring Databases
Weekly	Monitor Sybase disk usage		Monitoring and Tuning Databases
	Clean up old files		
Monthly	Reboot		Starting and Stopping Servers
Before and After Installations	Run DbVerify scripts		Installing Databases and Patches

Starting and Stopping Servers Procedures

- **Servers DBAs routinely start up and shut down include:**
  - **ASE Servers**
  - **ASE Backup Servers**
  - **ASE Monitor Server**
  - **SQS Servers**
  - **Replication Servers**

Database Devices & Logical Volumes
Database Devices

- **In order to create a new device, the DBA must have the following:**
  - **The name of database device to be created**
  - **A physical device on which to place database device**
  - **The device size in megabytes**
  - **The name of the mirror device, if one is in effect**

Database Devices & Logical Volumes
Database Device Procedure

```
/*****************************************************/
/* name: test_dev.sql                                */
/* purpose: allocate 3Mb device for testing          */
/* written: 12/18/97                                 */
/* revised:                                          */
/* reason:                                           */
/*****************************************************/
disk init name = test_dev,
physname                                              =
     "/usr/ecs/Rel_A/COTS/sybase/studentdevices/test_dev.dat
     ",
vdevno = 15,
size = 1536
go
sp_helpdevice test_dev
go
```

Installing Databases & Patches
ECS Assistant

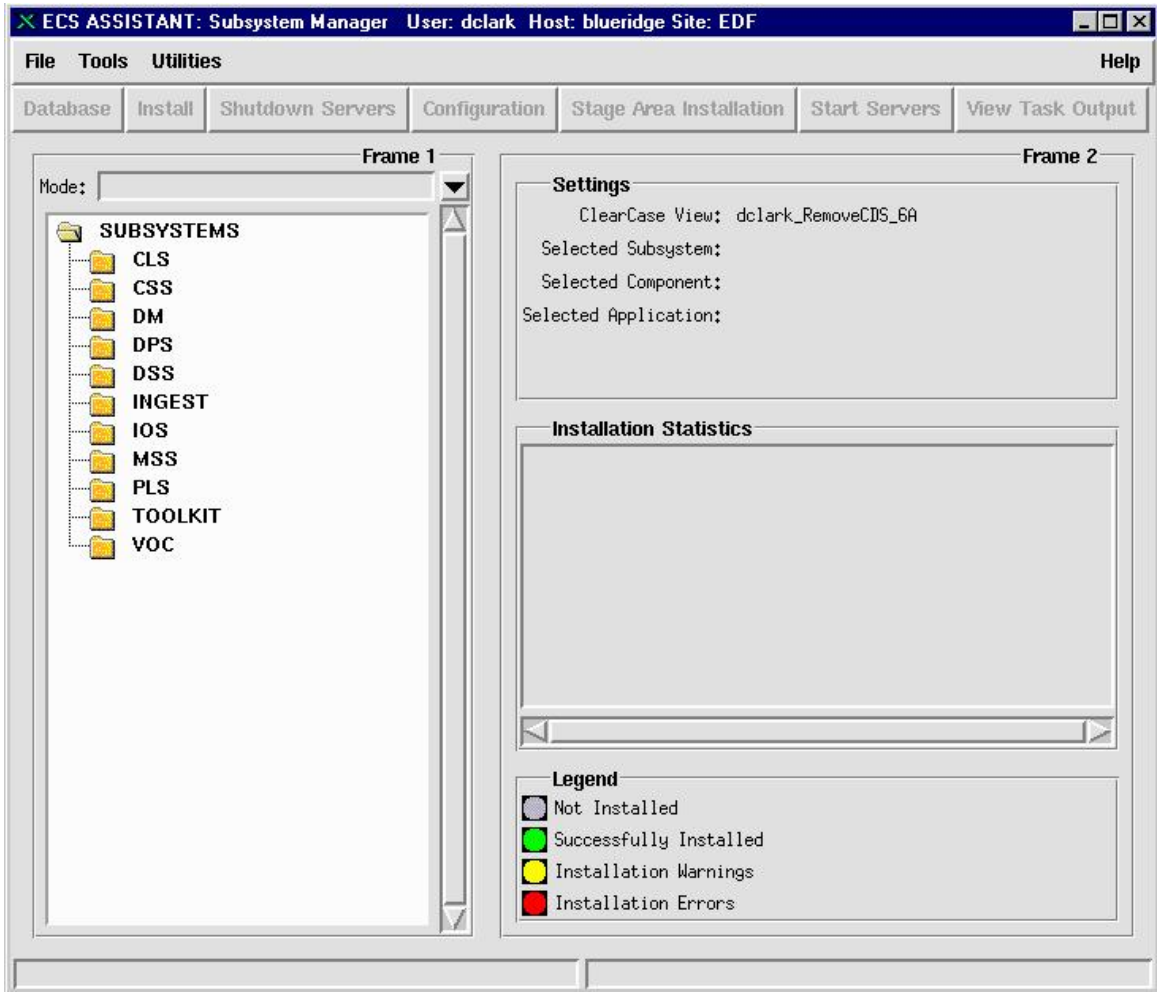

Configuring Databases
Configuration Parameter Tables

- **The configuration parameters are divided between two tables:**
  - **Sybase Configuration Parameter Table**
  - **DAAC-Specific Configuration Parameter Table**

Configuring Databases Procedures

- **Configuration parameters can be set or changed in one of two ways:**
  - **By executing the system procedure sp_configure with the appropriate parameters and values**
  - **By hand-editing your configuration file and then invoking sp_configure with the configuration file option**

Configuring Databases
sp_configure Sample Output

name	minimum	maximum	config value	run value
recovery interval	1	32767	0	5
allow updates	0	1	0	0
user connections	5	2147483647	0	25
memory	3850	2147483647	0	5120
open databases	5	2147483647	0	12
locks	5000	2147483647	0	5000
open objects	100	2147483647	0	500
procedure cache	1	99	0	20
fill factor	0	100	0	0
time slice	50	1000	0	100
database size	2	10000	0	2
tape retention	0	365	0	0
recovery flags	0	1	0	0
nested triggers	0	1	1	1
devices	4	256	0	10
remote access	0	1	1	1
remote logins	0	2147483647	0	20
remote sites	0	2147483647	0	10
remote connections	0	2147483647	0	20
pre-read packets	0	2147483647	0	3
upgrade version	0	2147483647	1002	1002
default sortorder id	0	255	50	50
default language	0	2147483647	0	0
language in cache	3	100	3	3
max online engines	1	32	1	1
min online engines	1	32	1	1
engine adjust interval	1	32	0	0
cpu flush	1	2147483647	200	200
i/o flush	1	2147483647	1000	1000
default character set id	0	255	1	1
stack size	20480	2147483647	0	28672
password expiration interval	0	32767	0	0
audit queue size	1	65535	100	100
additional netmem	0	2147483647	0	0
default network packet size	512	524288	0	512
maximum network packet size	512	524288	0	512
extent i/o buffers	0	2147483647	0	0
identity burning set factor	1	9999999	5000	5000
allow sendmsg	0	1	0	0
sendmsg starting port number	0	65535	0	0

Configuring Databases
Configuration Registry

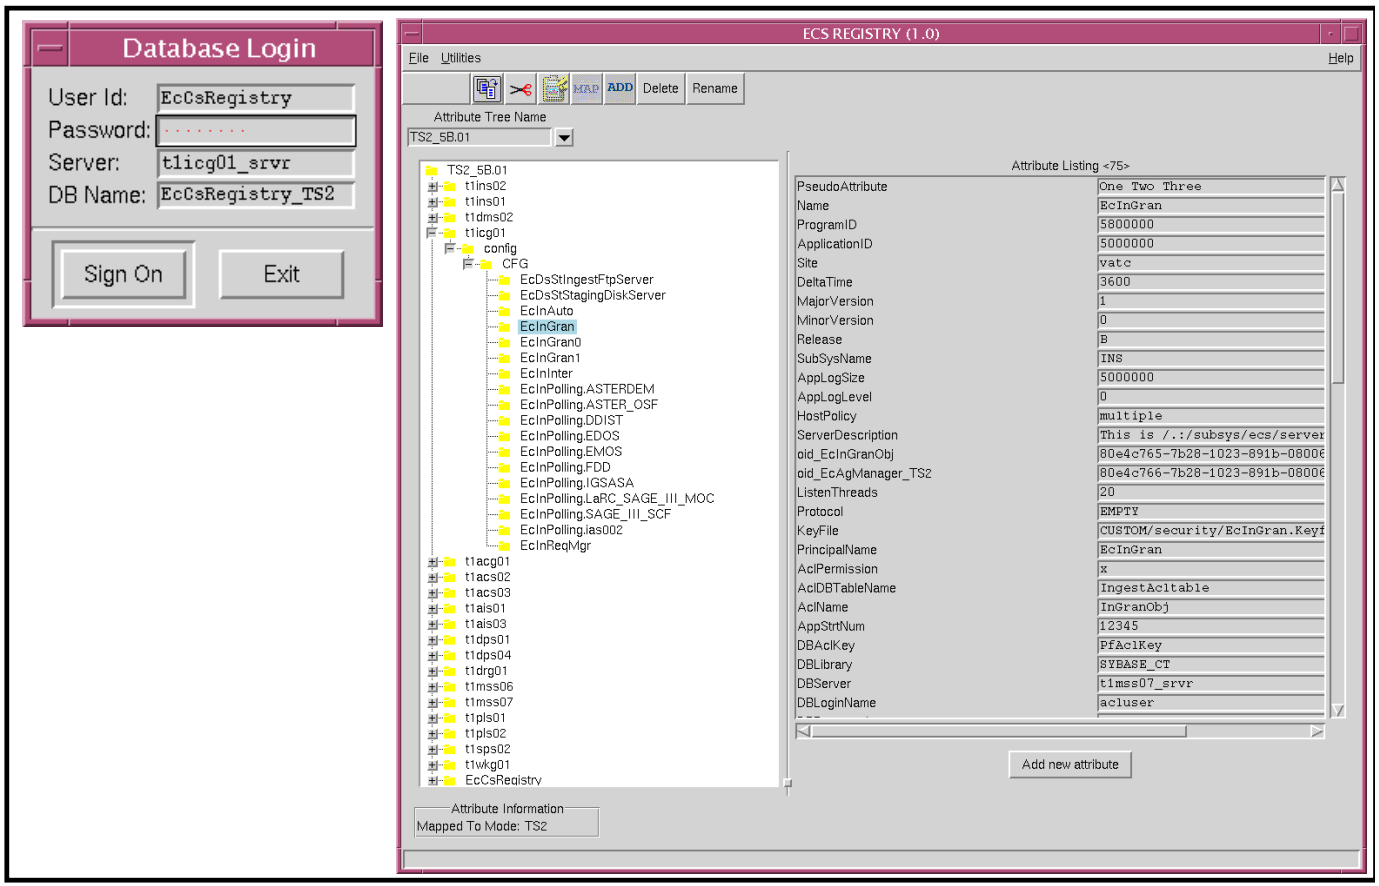

Indexes, Segments, & Caches
Indexes

- **Sybase allows the definition of two types of indexes:**
  - **Clustered index, where the rows in a database table are physically stored in sequence determined by the index.**
  - **Non-clustered indexes, which differ from their clustered counterpart in that the physical order of rows is not necessarily the same as their indexed order.**

Indexes, Segments, & Caches Segment Use

- **Segments are used when:**
  - **A table is placed on one device and its non-clustered indexes on a device on another disk controller, the time required to read or write to the disk can be reduced since disk head travel is usually reduced.**
  - **A large, heavily used table is split across devices on two separate disk controllers, read/write time may be improved.**
  - **The ASE Server stores the data for text and image columns on a separate chain of data pages. By default, this text chain is placed on the same segment as the table.**

Indexes, Segments, & Caches Segments

- **Subsystem databases, for example, consist of:**
  - **Default data segment used if no other segment specified in the create statement**
  - **SYSLOGS, transaction logs**
  - **System tables and indexes**
  - **OPS mode data segment**
  - **OPS mode index segment**
  - **TS1 mode data segment**
  - **TS1 mode index segment**
  - **TS2 mode data segment**
  - **TS2 mode index segment**

Indexes, Segments & Caches
Types of Caches

- **Default caches:**
  - **Data caches retain most recently accessed data and index pages**
  - **Procedure caches retain most recently accessed stored procedure pages**
  - **User transaction log caches are transaction log pages that have not yet been written to disk for each user**
- **Named caches, which are subdivisions of default caches**

Backing Up & Recovering Data Backups

- **Manual backups can be performed at any time by the DBA and are recommended for the following situations:**
  - **Any change to the master database, including new logins, devices, and databases**
  - **Any major change to user databases, such as a large ingest or deletion of data, definition of indexes**
  - **Other mission-critical activities as defined by the DAAC operations controller**
- **Automatic Backups**

Backing Up & Recovering Data

- **Performed when a database is corrupt or a device fails**
- **Run *dbcc* command frequently**
- **Dump current database and transaction log for failed database (or, if necessary, use most recent dumps)**
- **Set space defaults**
- **Drop database and device, and initialize new database device**
- **Re-create database**
- **Reload data from backups**
- **Manual recovery: System Administrator uses *load database* and *load transaction* commands**

Establishing Database Security
Discretionary Access Controls

Roles		Privileges
System Administrator	sa_role	Grant a specific user permissions needed to perform standard system administrator duties including: Installing ASE server and specific ASE server modules Managing the allocation of physical storage Tuning configuration parameters Creating databases
Site Security Officer	sso_role	Grant a specific user the permissions needed to maintain ASE server security including: • Adding server logins • Administrating passwords • Managing the audit system • Granting users all roles except the sa_role
Operator	oper_role	Grant a specific user the permissions needed to perform standard functions for the database including: • Dumping transactions and databases • Loading transactions and databases
Navigator	navigator_role	Grant a specific user the permissions needed to manage the navigation server
Replication	replication_role	Grant a specific user the permissions needed to manage the replication server
Sybase Technical Support	sybase_ts_role	Grant a specific user the permissions needed to execute database consistency checker (dbcc), a Sybase supplied utility supporting commands that are normally outside of the realm of routine system administrator activities

Establishing Database Security Identification & Authentication Controls

- **Providing users with access to servers and their databases consists of the following steps:**
  - **A server login account for a new user is created.**
  - **The user is added to a database and optionally assigned to a group.**
  - **The user or group is granted permissions on specific commands and database objects.**

Establishing Database Security
EMD Security Directive

- **All System Administrators and Database Administrators at the sites are responsible for reasonable security measures when installing custom software. This means:**
  - **Changing the permissions of online secure files to the minimum level required**
  - **Backing up secure file(s) to removable media (floppy or tape) and removal of secure files immediately after installation is complete and then keeping the removable medium in a secure location**

Copying, Replicating, & Extracting Data

- **Copy**
  - **Dump database to be copied**
  - **Create a database to load with the dump**
- **Bulk Copy (bcp)**
  - **Located in $SYBASE/OCS-12_5/bin directory**
  - **Need ASE Server account and appropriate permissions**
  - **Syntax:**

    bcp [[database_name].owner.]table_name {in | out} datafile [-e errfile] [-n] [-c] [-t field_terminator] [-r row_terminator] [-U username] [-S server]

Replication System Administration System Administrator Tasks

Task	Roles
Installing Replication Server	Replication System Administrator (RSA)
Adding or removing a Replication Server	RSA
Starting up and shutting down Replication Server.	RSA
Configuring Replication Server	RSA
Maintaining Routes (Creating and modifying)	RSA
Managing the RSSD	RSA
Adding a primary and replicate database.	RSA
Adding login names, database users, and administering appropriate permissions	RSA
Adding replicated tables or changing table schemas. Creating and modifying replicated tables Creating and modifying replication definitions Creating and materializing subscriptions at replicate sites.	RSA
Defining data server function-string classes and function strings.	RSA
Applying database recovery procedures.	RSA
Maintaining and monitoring database connections	RSA
Monitoring Replication Server	RSA
Processing rejected transactions	RSA
Quiescing Replication Server	RSA
Reconciling database inconsistencies.	RSA

Replication System Administration Database Administrator Tasks

Task	Roles
Installing Replication Server	DBA
Managing the RSSD	DBA
Adding a primary and replicate database.	DBA
Adding login names, database users, and administering appropriate permissions	DBA
Adding replicated tables or changing table schemas. Creating and modifying replicated tables Creating and modifying replication definitions Creating and materializing subscriptions at replicate sites.	DBA
Defining data server function-string classes and function strings.	DBA
Applying database recovery procedures.	DBA
Processing rejected transactions	DBA
Quiescing Replication Server	DBA
Reconciling database inconsistencies.	DBA

Performance Monitoring & Tuning Monitoring

- **Monitor database performance:**
  - **Before and after cache or pool configuration changes**
  - **Before and after certain sp_configure changes**
  - **Before and after the addition of new queries to your application mix**
  - **Before and after an increase or decrease in the number of Adaptive Server engines**
  - **When adding new disk devices and assigning objects to them**
  - **During peak periods, to look for contention**
  - **During stress tests to evaluate an Adaptive Server configuration for a maximum expected application load**
  - **When performance seems slow or the system behaves abnormally**

Performance Monitoring & Tuning
Tuning Options

Layers	Tuning Options
Application	Remote or replicated processing to move decision support off machine
	Stored procedures to reduce compilation time and network usage
	Minimum locking level that meets application needs
Database	Transaction log thresholds to automate dumps and avoid running out of space
	Thresholds for space monitoring in data segments
	Partitions to speed loading of data
	Devices to avoid disk contention, take advantage of I/O parallelism
Server	Tuning memory, most critical configuration parameters and other parameters
	Configuring cache sized and I/O sizes
	Scheduling batch jobs and reporting for off hours
	Reconfiguring parameters for shifting workload patterns
Devices	More medium-sized devices and more controllers for better I/O throughput
	Distributing databases, tables, and indexes for even I/O load across devices
	Segments, partitions for I/O performance on large tables used for parallel queries
Network	Configuring packet sizes to match application needs
	Configuring subnets
	Isolating heavy network uses
	Configuring for multiple network engines
Hardware	Configuring the housekeeper task to improve CPU use
	Configuring multiple data caches
Operating System	Choosing between riles and raw partitions
	Increasing memory size

Ensuring Database Quality Integrity Monitoring

- **The integrity of the internal structures of a database depends upon the System Administrator or Database Owner running database consistency checks on a regular basis. Two major functions of dbcc are:**
  - **Checking allocation structures (the commands checkalloc, tablealloc, and indexalloc).**
  - **Checking page linkage and data pointers at both the page level and row level (checktable and checkdb).**

Sybase Troubleshooting
Space Usage

- **Thresholds are defined on segments to provide a free space value at which a procedure is executed to provide a warning or to take remedial action.**
- **Use sp_addthreshold to define your own thresholds:**
  - **sp_addthreshold database_name, segment_name, free_space, procedure_name**
    - **free_space is the number of free pages at which the threshold procedure executes**
    - **procedure_name is the stored procedure which the threshold manager executes when the number of free pages falls below the free_space value**

Troubleshooting Deadlocks

- A deadlock (also known as a "deadly embrace") is a situation where two database processes are simultaneously attempting to lock data that the other holds
  - For example, two users (A and B) are updating the same table of data at the same time
    - User A holds a lock on Page 1 and requests a lock on Page 2
    - Meanwhile, user B holds a lock on Page 2 and has requested a lock on Page 1
    - Without intervention, these two jobs would never finish

Oracle
Basic Procedures

- **Basic Oracle procedures required to run PDS include:**
  - **Starting up the database**
  - **Shutting down the database**
  - **Controlling the listener**
  - **Using the data dictionary**
  - **Obtaining archiving information**
  - **Obtaining group and member information**
  - **Controlling log switches and checkpoints**
  - **Troubleshooting**
  - **Accessing a dynamic performance view**
  - **Displaying parameter values**
  - **Displaying information about users**
  - **Displaying system and object privilege information**